




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : H04N 7/26, 7/50, 7/32	A1	(11) International Publication Number: WO 00/46996 (43) International Publication Date: 10 August 2000 (10.08.00)
(21) International Application Number: PCT/US00/02800 (22) International Filing Date: 3 February 2000 (03.02.00) (30) Priority Data: 60/118,359 3 February 1999 (03.02.99) US 09/388,653 2 September 1999 (02.09.99) US (71) Applicant: SARNOFF CORPORATION [US/US]; 201 Washington Road, CN5300, Princeton, NJ 08543-5300 (US). (72) Inventors: SETHURAMAN, Sriram; Kensington Arms Apartment 28B, Hightstown, NJ 08520 (US). KRISHNA-MURTHY, Ravi; 5614 Hunters Glen Drive, Plainsboro, NJ 08536 (US). (74) Agents: NEY, Andrew, L.; Ratner & Prestia, 301, One Westlakes (Berwyn), P.O. Box 980, Valley Forge, PA 19482-0980 (US) et al.		(81) Designated States: BR, CA, CN, JP, KR, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: FRAME-LEVEL RATE CONTROL FOR PLUG-IN VIDEO CODECS



(57) Abstract

A variety of different types of video frame encoders can be configured with, e.g., a multimedia processing subsystem, as long as the video frame encoder conforms to the interface protocol of the subsystem. A video controller in the subsystem performs the higher-level functions of coordinating the encoding of the video stream, thereby allowing the video frame encoder to limit its processing to the lower, frame level. In particular, the video controller provides information needed by the video frame encoder to encode the current frame in the video sequence. In addition to the raw image data, this information includes the type of frame to be encoded (e.g., an I or P frame), the currently available bandwidth for encoding the current frame, the time since the previous encoded frame, the desired frame rate, and a quality measure that may be used to trade off spatial and temporal qualities. The video frame encoder either encodes the frame as requested or indicates to the video controller that the frame should be skipped or otherwise not encoded as requested. The video controller can then respond appropriately, e.g., by requesting the video frame encoder to encode the next frame in the video sequence.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon	KR	Republic of Korea	PL	Poland		
CN	China	KZ	Kazakhstan	PT	Portugal		
CU	Cuba	LC	Saint Lucia	RO	Romania		
CZ	Czech Republic	LI	Liechtenstein	RU	Russian Federation		
DE	Germany	LK	Sri Lanka	SD	Sudan		
DK	Denmark	LR	Liberia	SE	Sweden		
EE	Estonia			SG	Singapore		

FRAME-LEVEL RATE CONTROL FOR PLUG-IN VIDEO CODECS
BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to image processing, and, in particular, to video
5 compression.

Cross-Reference to Related Applications

This application claims the benefit of the filing date of U.S. provisional application
no. 60/118,359, filed on 02/03/99 as attorney docket no. SAR 13331P.

Description of the Related Art

The goal of video compression processing is to encode image data to reduce the
number of bits used to represent a sequence of video images while maintaining an acceptable
level of quality in the decoded video sequence. This goal is particularly important in certain
15 applications, such as videophone or video conferencing over POTS (plain old telephone
service) or ISDN (integrated services digital network) lines, where the existence of limited
transmission bandwidth requires careful control over the bit rate, that is, the number of bits
used to encode each image in the video sequence. Furthermore, in order to satisfy the
transmission and other processing requirements of a video conferencing system, it is often
20 desirable to have a relatively steady flow of bits in the encoded video bitstream. That is, the
variations in bit rate from image to image within a video sequence should be kept as low as
practicable.

Achieving a relatively uniform bit rate can be very difficult, especially for video
compression algorithms that encode different images within a video sequence using different
25 compression techniques. Depending on the video compression algorithm, images may be
designated as the following different types of frames for compression processing:

- o An intra (I) frame which is encoded using only intra-frame compression techniques,
- o A predicted (P) frame which is encoded using inter-frame compression techniques
based on a previous I or P frame, and which can itself be used as a reference frame to encode
30 one or more other frames,

- o A bi-directional (B) frame which is encoded using bi-directional inter-frame compression techniques based on a previous I or P frame, a subsequent I or P frame, or a combination of both, and which cannot itself be used to encode another frame, and
- o A PB frame which corresponds to two images -- a P frame and a subsequent B frame -- that are encoded as a single frame (as in the H.263 video compression algorithm).

Depending on the actual image data to be encoded, these different types of frames typically require different numbers of bits to encode. For example, I frames typically require the greatest number of bits, while B frames typically require the least number of bits.

In a typical transform-based video compression algorithm, a block-based transform, such as a discrete cosine transform (DCT), is applied to blocks of image data corresponding either to pixel values or pixel differences generated, for example, based on a motion-compensated inter-frame differencing scheme. The resulting transform coefficients for each block are then quantized for subsequent encoding (e.g., run-length encoding followed by variable-length encoding). The degree to which the transform coefficients are quantized directly affects both the number of bits used to represent the image data and the quality of the resulting decoded image. This degree of quantization is also referred to as the quantization level, which is often represented by a specified quantizer value that is used to quantize all of the transform coefficients. In some video compression algorithms, the quantization level refers to a particular table of quantizer values that are used to quantize the different transform coefficients, where each transform coefficient has its own corresponding quantizer value in the table. In general, higher quantizer values imply more severe quantization and therefore fewer bits in the encoded bitstream at the cost of lower playback quality of the decoded images. As such, the quantizer is often used as the primary variable for controlling the tradeoff between bit rate and image quality.

At times, using quantization level alone may be insufficient to meet the bandwidth and quality requirements of a particular application. In such circumstances, it may become necessary to employ more drastic techniques, such as frame skipping, in which one or more frames are dropped from the video sequence. Such frame skipping may be used to sacrifice short-term temporal quality in the decoded video stream in order to maintain a longer-term spatial quality at an acceptable level.

SUMMARY OF THE INVENTION

The present invention is directed to video encoding techniques that separate the functionality for controlling the higher-level (i.e., sequence-level) aspects of encoding video data from the functionality for implementing the lower-level (i.e., frame-level) encoding of individual video frames within the video sequence. The techniques of the present invention enable video processing systems to be build modularly, where a video processing subsystem that controls the sequence-level processing can be configured with any of a variety of plug-in video encoders that control the frame-level processing that conform to the interface protocol of the subsystem. This enables the selection of video encoder to be dependent on the particular application. For example, more expensive, higher-quality video encoders can be employed for higher-quality applications, while less expensive, lower-quality video encoders can be employed for lower-quality applications.

The present invention allows control parameters such as bit rate, desired spatio-temporal quality, and key-frame requests to be set at any or every frame over a video sequence, thus allowing the encoding to be tailored dynamically to network conditions, user preferences, and random access/re-synchronization requirements.

In one embodiment, the present invention is a method for encoding a video sequence by a video encoder, comprising the steps of (a) receiving a current frame of video data; (b) receiving a set of input parameter values corresponding to the current frame; (c) determining whether to skip the current frame based on the set of input parameter values; (d) if appropriate, encoding the current frame based on the set of input parameter values; and (e) repeating steps (a)-(d) for one or more other frames in the video sequence, wherein one or more of the input parameter values varies from frame to frame in the video sequence.

BRIEF DESCRIPTION OF THE DRAWINGS

Other aspects, features, and advantages of the present invention will become more fully apparent from the following detailed description, the appended claims, and the accompanying drawings in which:

Fig. 1 shows a multimedia processing system, according to one embodiment of the present invention;

Fig. 2 provides pseudocode for the function rcInitGOP;

Fig. 3 shows a flow diagram of the processing performed by the function rcInitGOP of Fig. 2;

Fig. 4 provides pseudocode for the function rcFrameSkip;

Fig. 5 shows a flow diagram of the processing performed by the function rcFrameSkip of Fig. 4;

Fig. 6 provides pseudocode for the function rcGetTarget; and

Fig. 7 shows a flow diagram of the processing performed by the function rcGetTarget of Fig. 6.

DETAILED DESCRIPTION

Fig. 1 shows a multimedia processing system 100, according to one embodiment of the present invention. System 100 encodes sequences of video images and optionally combines the resulting compressed video bitstream with audio and/or data streams for transmission over a network 102. System 100 may be used for multimedia applications such as broadcasting or teleconferencing.

In particular, frame acquisition module 104 (e.g., a video camera) provides a sequence of video images to video controller 106, which passes the frames to video frame encoder 108, which performs the actual encoding of the video data. The resulting encoded video bitstream 118 is transmitted by encoder 108 to buffer and multiplex module 110, which may combine the encoded video bitstream with one or more audio and/or data streams provided by audio/data module 112 to form a multimedia stream 120, which is then transmitted over network 102 to one or more remote client receivers (not shown) that decode the multimedia stream for playback.

In system 100, video controller 106 controls the operations of all of the other modules in system 100. In particular, video controller 106 controls the operations of frame acquisition module 104, buffer and multiplex module 110, and audio/data module 112. Video controller 106 also does all the multiplexing and handshaking with the remote clients. As a result, encoder 108 can be implemented independent of video controller 106, with an agreed-upon interface 116 between the controller and the encoder.

In general, the video controller need not control all modules in the system. System 100 is provided to show an example of the separation of a system-level application from the plug-in video codec. Other system configurations are possible. For example, audio can have

its own controller. The arrows from video controller 106 are shown mainly to indicate the communication needed to ensure proper synchronization of video with audio and data.

The independent implementation of video frame encoder 108 is indicated in Fig. 1 by depicting encoder 108 outside of a multimedia processing subsystem 114 that contains video controller 106 and the other modules 104, 110, and 112. As long as a video frame encoder conforms to the protocol of interface 116, it can be configured with subsystem 114 to form a multimedia processing system similar to that shown in Fig. 1. This modularity enables subsystem 114 to be configured with different video frame encoders, depending on the particular application. For lower quality applications, subsystem 114 can be configured with lower quality, less expensive video frame encoders. Similarly, for higher quality applications, subsystem 114 can be configured with higher quality, more expensive video frame encoders. Moreover, as technology improves and better video frame encoders become available, a multimedia processing system can be upgraded by reconfiguring subsystem 114 with a newer, better video frame encoder. Of course, in each configuration, a matched video frame decoder is used at each client end of the network.

In general, in order for video frame encoder 108 to operate as independent module, video controller 106 provides video frame encoder 108 all of the information needed to compress a sequence of video frames into an encoded video bitstream. This information includes certain parameters used by encoder 108 to control the video compression processing as well as the raw video data itself.

The present invention addresses the problem of frame-level rate control at video frame encoder 108. The rate control algorithm of the present invention, implemented at the encoder, reacts to the video controller's demands. In addition, the rate control algorithm proactively allocates bandwidth based on scene content and "slack," where "slack" is the state of a virtual buffer that is maintained at the video frame encoder. The rate of depletion of this virtual buffer is calculated from the parameters sent from the video controller.

The rate control algorithm of the present invention takes much of the responsibility of frame-level rate control from the video controller and moves it to the frame encoder. This is a very desirable property since it enables the encoder to proactively allocate bits based on scene content, instead of just reacting to the controller parameters. Also, very little rate control has to be performed by the controller. However, due to fluctuations in the video

data being encoded, it may sometimes be necessary for the controller to perform simple corrections, like skipping a frame.

In preferred embodiments, the video controller passes most or all of the following information to the video frame encoder when asking for the encoding of a frame:

5 o **bitrate**: The instantaneous available bandwidth (bit rate) which the encoder is expected to derive. The value of bitrate may be based on feedback received by controller 106 from buffer module 110 and network 102.

 o **I_frame**: Flag indicating whether or not the current frame is a key frame (i.e., an intra-coded picture).

10 o **Quality**: Degree of spatial vs. temporal quality that the controller desires. The value of Quality, which can change from frame to frame, can range from 0 to 100, where 0 corresponds to maximal temporal quality and 100 corresponds to maximal spatial quality. Maximal temporal quality means that a uniform frame rate is preserved even if severe quantization is needed to achieve target bit rates. Maximal spatial quality means that frames
15 can be freely dropped if appropriate to maintain a specified degree of spatial quality within the coded frames. In between these two extremes, the Quality value corresponds to monotonically decreasing temporal resolution and corresponding spatial qualities that are content dependent.

20 o **skip_time_in_ms**: Time from previous encoded frame, which indicates number of frames skipped from previous encoded frame.

 o **target_frame_rate**: Desired frame rate at which the controller wants the encoder to operate.

Additional parameters relating to buffer state (like fullness) may also be passed, for finer control. Also, during an initialization phase, the following static parameters are passed:

25 o **Frame size** (i.e., width and height);

 o **Input format**; and

 o **Periodicity of key frames** (e.g., after every 50 coded frames, there will be an I-frame).

This is not required for the case where frames passed by the controller can be skipped by the encoder.

30 The following discussion assumes only I and P frames with variables defined as follows:

- o `def_frame_rate`: Default frame rate.
- o `first_frame_in_sequence`: Flag indicating whether the current I frame is the first frame in the video sequence. If so, then certain parameter values (e.g., `prev_bitrate`, `prev_bitCount`, `prev_I_frame_bits`, and `QPIprev`) will not be available.
- 5 o `frame_rate`: Number of frames per second.
- o `freeze_time`: Current time that is available for I-frame transmission.
- o `I_SENSITIVITY`: Controls how the Quality parameter affects the QP for I frames.
- o `MAX_BUFDELAY_CONST`: Maximum buffer delay that can be tolerated based on overall system delay requirements. This value can be expressed as a constant times
- 10 `P_frame_time`.
- o `MAX_FREEZE_TIME`: Maximum time that is available for I-frame transmission. During this time, the previous frame will be held in the decoder-side display buffer. It is desirable to keep this freeze time within certain bounds (e.g., 300 ms).
- o `MIN_BITS_REQUIRED_FOR_CODING`: Minimum number of bits required for
- 15 coding the current frame.
- o `P_frame_time`: Average time duration between P frames at a given time. This value could vary depending on the current load on the encoder, the particular algorithm being used, and the capture delays.
- o `prev_bitCount`: Number of bits actually used to encode the previous frame (whether
- 20 I or P).
- o `prev_bitrate`: Target number of bits for encoding the previous frame (whether I or P).
- o `prev_I_frame_bits`: Number of bits used to encode the previous I frame.
- o `P_SENSITIVITY`: Controls how the Quality parameter affects the QP for P frames.
- 25 o `QP_DEFAULT`: Default value for the quantizer parameter QP that is used for an I frame when no information is available.
- o `QPI`: Quantizer parameter for current I frame.
- o `QPIprev`: Quantizer parameter for previous I frame.
- o `Rslack`: Number of bits left in the encoder virtual buffer.
- 30 o `R_WINDOW`: The time window over which slack can be accumulated and distributed. `R_WINDOW` is related to buffer size. For low-delay applications,

R_WINDOW is very small. For streaming or broadcast applications, R_WINDOW can be large.

- o target_bits: Target number of bits for encoding the current frame.

5 Using these terms, the following rate control algorithm is implemented, which constantly reacts to the changing parameters supplied by the controller.

Rate Control For a Key Frame

10 At a key frame (i.e., an I frame at the beginning of a group of pictures (GOP)), the video controller calls a function rcInitGOP into the video frame encoder to calculate the target bit rate (target_bits) and set the initial value for the quantizer parameter QPI for the I frame. The function rcInitGOP receives the following input parameters: bitrate, Quality, skip_time_in_ms, prev_bitCount, and first_frame_in_sequence. Pseudocode for the function rcInitGOP is given in Fig. 2.

15 Fig. 3 shows a flow diagram of the processing performed by the function rcInitGOP of Fig. 2. If the current I frame is the first frame in the video sequence (step 302 and line 5 of Fig. 2), then there are no values for certain parameters related to previous frames and special processing is implemented. In particular, the time available for transmission of this I frame (freeze_time) is generated based on the maximum freeze time
20 (MAX_FREEZE_TIME) and the quality parameters (Quality and I_SENSITIVITY) (step 304 and line 7). This current freeze time is then used to generate the target for the current frame (target_bits) (step 306 and line 8) and the quantizer parameter for the current I frame (QPI) is generated from the default quantizer parameter value (QP_DEFAULT) and the quality parameters (Quality and I_SENSITIVITY) (step 308 and line 10)

25 When the current I frame is not the first frame in the video sequence, the slack in the buffer (Rslack) at the start of the new GOP is generated (step 310 and line 14). If the buffer slack is too large (step 312 and line 20), then the current frame is not encoded as an I frame (step 314 and line 20). In this case, the video controller is informed by the video frame encoder that the current frame was not encoded as an I frame, and the video controller will
30 respond with some appropriate action (e.g., request the video frame encoder to encode the next frame as an I frame).

If the buffer slack is not too large, then the time available for transmission of this I frame (freeze_time) is generated based on the maximum freeze time (MAX_FREEZE_TIME) and the quality parameters (Quality and I_SENSITIVITY) (step 316 and line 24). This current freeze time is then used to generate the target for the current frame (target_bits) (step 318 and line 27) and the quantizer parameter for the current I frame (QPI) is generated from the previous quantizer parameter value (QPIprev), the number of bits used to encode the previous I frame (prev_I_frame_bits) and the current target (target_bits) (step 320 and line 30).

Whether or not the current I frame is the first frame in the sequence, the quantizer parameter for the current frame (QPI) is clipped to the allowable range (e.g., 1 to 31) (step 322 and line 33).

The rest of the processing shown in Fig. 3 occurs after the function rcInitGOP has completed. The value for QPI returned from the function rcInitGOP is used first to encode the I frame with no quantizer adaptation over the I frame (step 324). If the number of bits actually used to encode the I frame sufficiently match the target (i.e., target_bits) computed in by the function rcInitGOP (i.e., to within a specified tolerance) (step 326), then processing of the I frame is complete. If, on the other hand, the actual number of bits does not sufficiently match the target bit rate, then the value of QPI is recalculated using a linear model (step 328) based on the following formula:

$$QPI(new) = (QPI(old) * actual_bits_with_QPI(old)) / target_bits$$

and the I frame is re-encoded using the new QPI value QPI_new (steps 322 and 324). This recalculation of QPI and re-encoding of the I frame may be repeated one or more times until the target bit rate is achieved (or until the linear model generates a value for QPI that has already been used (not shown in Fig. 3)).

Rate Control for a P frame

For P frames, the video controller calls a function rcFrameSkip into the video frame encoder to decide whether or not to skip the current frame based on the current bit rate (bitrate), the time between the previous encoded frame and the current frame (skip_time_in_ms), and the actual number of bits used to encode the previous frame (prev_bitCount). Pseudocode for the function rcFrameSkip is given in Fig. 4.

Fig. 5 shows a flow diagram of the processing performed by the function rcFrameSkip of Fig. 4. The current slack in the buffer (Rslack) is generated (step 502 and line 3 of Fig. 4), and then clipped if it is too low (step 504 and line 9). An average time used to encode a P frame is then updated (step 506 and line 15). The update weights the current frame more than the past history, to reflect changing dynamics of processor performance and buffer status. If the buffer slack is too large (step 508 and line 16), then the current frame is to be skipped (i.e., not encoded) (step 510 and line 16). The video controller is informed by the video frame encoder that the current frame is to be skipped, and the video controller will respond with some appropriate action (e.g., request the video frame encoder to encode a subsequent frame (e.g., the next frame) as a P or I frame, depending on the location of the current frame in the GOP). Otherwise, the current frame is to be encoded as a P frame (step 512 and line 17).

If the function rcFrameSkip indicates that the current frame is to be encoded as a P frame, then the video controller calls a function rcGetTarget into the video frame encoder to generate the target number of bits used to encode the current P frame, based on the quality parameter (Quality), the current bit rate (bitrate), and the minimum number of bits to be used to encode this frame (MIN_BITS_REQUIRED_FOR_CODING). Pseudocode for the function rcGetTarget is given in Fig. 6.

Fig. 7 shows a flow diagram of the processing performed by the function rcGetTarget of Fig. 6. An execution-constrained frame rate (frame_rate) is generated (step 702 and line 5 of Fig. 6), and a quality-constrained frame rate (target_frame_rate) is generated (step 704 and line 6). If the quality-constrained frame rate is greater than the execution-constrained frame rate, then the quality-constrained frame rate is limited to the execution-constrained frame rate (step 706 and line 9).

The target number of bits to encode the current frame is then generated (step 708 and line 10), based on the quality-constrained frame rate (target_frame_rate) and the current buffer state (Rslack). If the target number of bits is too small (step 710 and line 13), then the current frame is to be skipped (step 712 and line 13). The video controller is informed by the video frame encoder that the current frame is to be skipped, and the video controller will respond with some appropriate action (e.g., request the video frame encoder to encode a subsequent frame (e.g., the next frame) as a P or I frame, depending on the location of the

current frame in the GOP). Otherwise, the current frame is to be encoded as a P frame and the target number of bits is returned by the function rcGetTarget (step 714 and line 14).

After successfully generating a target number of bits to be used to encode the current P frame, the frame may be encoded using any suitable compression algorithm, including
5 those in which the quantizer parameter is allowed to vary from macroblock to macroblock within each image. As in the case of I frames, depending on the application, each P frame may be encoded one or more times to ensure acceptable achievement of the target bit rate.

Variations

10 In addition to deciding to skip frames based on buffer state and/or minimum bits to encode, the decision as to whether to skip a frame may be made after motion estimation. For example, if the number of bits required to encode just the motion vectors themselves is close to the target, then the frame may be skipped.

In some applications, the exact buffer information may be available. In those cases,
15 the calculations of Rslack can be replaced by using the exact buffer fullness values.

In the discussion above, the video frame encoder indicates to the video controller whether the current frame is to be skipped, and the video controller handles the actual skipping of frames (e.g., selecting the next frame to encode). In some applications, this functionality may not be supported, and the video frame encoder may be required to encode
20 each frame provided by the video controller. One possibility is for the video frame encoder to encode the frame as a P frame with each and every macroblock designed as a skipped block. This will effectively achieve the same result of skipping the frame with only a relatively small amount of overhead data to be transmitted.

The frame-level rate control of the present invention is particularly applicable in real-
25 time very-low-bit-rate coders. The algorithm can be implemented by coders using only P frames as well as coders using PB-frames, an added functionality when compared to the TMN8 test model for the H.263+ standard. The algorithm provides a buffer delay variable, which can be selected by a user to trade delay for graceful change in spatial quality over time. By adapting the target bit allocation for a frame to the scene content, spatial quality is
30 maintained in high-motion areas and quick recovery is made possible after an abrupt motion, while maintaining the buffer delay within desired limits.

The present invention can be embodied in the form of methods and apparatuses for

practicing those methods. The present invention can also be embodied in the form of program code embodied in tangible media, such as floppy diskettes, CD-ROMs, hard drives, or any other machine-readable storage medium, wherein, when the program code is loaded into and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the invention. The present invention can also be embodied in the form of program code, for example, whether stored in a storage medium, loaded into and/or executed by a machine, or transmitted over some transmission medium, such as over electrical wiring or cabling, through fiber optics, or via electromagnetic radiation, wherein, when the program code is loaded into and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the invention. When implemented on a general-purpose processor, the program code segments combine with the processor to provide a unique device that operates analogously to specific logic circuits.

It will be further understood that various changes in the details, materials, and arrangements of the parts which have been described and illustrated in order to explain the nature of this invention may be made by those skilled in the art without departing from the principle and scope of the invention as expressed in the following claims.

What is claimed is:

1. A method for encoding a video sequence by a video encoder, comprising the steps of:
 - (a) receiving a current frame of video data;
 - (b) receiving a set of input parameter values corresponding to the current frame;
 - 5 (c) determining whether to skip the current frame based on the set of input parameter values; (d) if appropriate, encoding the current frame based on the set of input parameter values; and
 - (e) repeating steps (a)-(d) for one or more other frames in the video sequence, wherein one or more of the input parameter values vary from frame to frame in the video sequence.
- 10 2. The invention of claim 1, wherein step (d) comprises the steps of:
 - (1) generating a target bit count for the current frame;
 - (2) selecting one or more quantization parameter (QP) values for the current frame based on the target bit count; and
 - (3) encoding the current frame based on the one or more QP values.
- 15 3. The invention of claim 2, wherein step (d) further comprises the steps of:
 - (4) determining whether an actual number of bits used to encode the current frame sufficiently matches the target bit count;
 - (5) if appropriate, adjusting at least one of the one or more QP values based on the actual number of bits and the target bit count; and
 - 20 (6) re-encoding the current frame based on the adjusted one or more QP values.
4. The invention of claim 1, wherein:

the set of input parameter values comprises a quality parameter used to trade off spatial quality versus temporal quality; and

step (d) comprises the step of encoding the current frame based on the quality parameter.
- 25 5. The invention of claim 4, wherein step (d) comprises the steps of:
 - (1) generating a target bit count for the current frame based on the quality parameter; and
 - (2) encoding the current frame based on the target bit count.

6. The invention of claim 5, wherein the target bit count is generated based on the quality parameter, a current bit rate value, and a buffer slack measure.

7. The invention of claim 5, wherein, when the current frame is an I frame:

$$\text{freeze_time} = \text{MAX_FREEZE_TIME} * (1.0 + ((\text{float}) (\text{Quality} - 50)) / \text{I_SENSITIVITY})$$

and

$$\text{target_bits} = (\text{freeze_time} * \text{bitrate}) / 1000;$$

wherein:

freeze_time is a current time that is available for I-frame transmission;

MAX_FREEZE_TIME is a specified maximum time that is available for I-frame transmission;

Quality is the quality parameter;

I_SENSITIVITY is a sensitivity parameter for I frames;

target_bits is a target number of bits for encoding the current frame; and

bitrate is the current bit rate value.

8. The invention of claim 5, wherein, when the current frame is a P frame:

$$\text{target_frame_rate} = \text{def_frame_rate} * (1.0 - (\text{Quality} - 50) / \text{P_SENSITIVITY});$$

and

$$\text{target_bits} = (\text{bitrate} / \text{target_frame_rate}) - ((\text{Rslack} * \text{P_frame_time}) / (\text{R_WINDOW}));$$

wherein:

target_frame_rate is a specified desired frame rate for the video encoder;

def_frame_rate is a specified default frame rate for the video encoder;

Quality is the quality parameter;

P_SENSITIVITY is a sensitivity parameter for P frames; and

target_bits is a target number of bits for encoding the current frame; and

bitrate is the current bit rate value;

Rslack is a number of bits left in a virtual buffer of the video encoder;

P_frame_time is a current average time duration between P frames; and

R_WINDOW is a time window over which buffer slack can be accumulated and distributed.

9. The invention of claim 4, further comprising the step of disabling the quality parameter for P frame target allocation, when an instantaneous execution-constrained frame rate is less than a frame rate generated from the quality parameter.

10. The invention of claim 1, wherein:

5 the set of input parameter values comprises a time since a previous encoded frame; and step (c) comprises the step of determining whether to skip the current frame based on the time since the previous encoded frame.

11. The invention of claim 10, wherein step (c) comprises the steps of:

10 (1) updating a measure of buffer fullness based on the time since the previous encoded frame; and

(2) determining whether to skip the current frame based on the measure of buffer fullness

12. The invention of claim 10, further comprising the step of using the time since the previous encoded frame to generate an instantaneous execution-constrained frame rate.

13. The invention of claim 1, wherein:

15 the set of input parameter values comprises a current bit rate value; and step (d) comprises the step of encoding the current frame based on the current bit rate value to enable the video encoder to provide variable bit-rate encoding.

14. The invention of claim 1, wherein, when the current frame is an I frame, step (c) comprises the steps of determining whether to skip the current frame, in which case a next
20 encoded frame is encoded as an I frame.

15. The invention of claim 1, wherein, when the current frame is a P frame, step (c) comprises the steps of determining whether to skip the current frame, in which case each macroblock in the current P frame is encoded as a skipped macroblock.

16. The invention of claim 1, wherein, when the current frame is a P frame, step (c)
25 comprises the steps of determining whether to skip the current frame after performing motion estimation based on a number of bits required to encode motion vectors from the motion estimation.

17. A video encoder for encoding a video sequence, comprising:

(a) means for receiving a current frame of video data;

(b) means for receiving a set of input parameter values corresponding to the current frame;

5 (c) means for determining whether to skip the current frame based on the set of input parameter values;

(d) means for encoding the current frame based on the set of input parameter values, wherein one or more of the input parameter values vary from frame to frame in the video sequence.

10 18. A machine-readable medium, having encoded thereon program code, wherein, when the program code is executed by a machine, the machine implements the steps of:



(a) receiving a current frame of video data in a video sequence;

(b) receiving a set of input parameter values corresponding to the current frame;

15 (c) determining whether to skip the current frame based on the set of input parameter values; (d) if appropriate, encoding the current frame based on the set of input parameter values; and

(e) repeating steps (a)-(d) for one or more other frames in the video sequence, wherein one or more of the input parameter values vary from frame to frame in the video sequence.

1/6




2/6

```

1 int rclnitGOP (int bitrate, int Quality, int skip_time_in_ms, int prev_bitCount,
  int first_frame_in_sequence)
2 {
3     int QPI, freeze_time;
4     prev_bitrate = bitrate;
5     if (first_frame_in_sequence) { /* prev_I_frame_bits and QPIprev are not
        available, */
6         /* Modify the freeze time based on quality. */
7         freeze_time = MAX_FREEZE_TIME* (1.0 + ((float) (Quality-50))
            /I_SENSITIVITY);
8         target_bits = (freeze_time*bitrate) /1000;
9         /* Set QPI modified by quality. */
10        QPI = (int) QP_DEFAULT*(1.0 - ((float) (Quality-50)) /I_SENSITIVITY);
11    }
12    else{
13        /* Compute slack in buffer at start of new GOP. */
14        Rslack += prev_bitCount - (prev_bitrate*skip_time_in_ms) /1000;
15        /* In slack calculations, the instantaneous bit rate at the previous encoded
            frame is held */
16        /* until the current encoded frame. */
17
18        /* If slack is to large, do not encode this frame. */
19        /* Next frame is encoded as key frame, if buffer constraints allow. */
20        if (Rslack > (bitrate* (MAX_BUFDELAY_CONST*P_frame_time))
            /1000) return -1;
21
22        /* Find a QP for the I-frame, based on freeze time and user selected
            quality. */
23        /* Modify freeze time according to user selected quality. */
24        freeze_time = MAX_FREEZE_TIME* (1.0 + ((float) (Quality-50))
            /I_SENSITIVITY);
25
26        /* Compute target bits for the modified freeze time. */
27        target_bits = (freeze_time*bitrate) /1000;
28
29        /* Based on last I-frame's QP and bitcount, obtain QP for current I-frame. */
30        QPI = (prev_I_frame_bits*QPIprev) /target_bits;
31    }
32    /* Clip QPI to [1,31] for the common standards. */
33    QPI = (QPI > 31)?31: ((QPI < 1)?1:QPI);
34    return QPI;
35 }

```

FIG. 2

**FIG. 3**

4/6

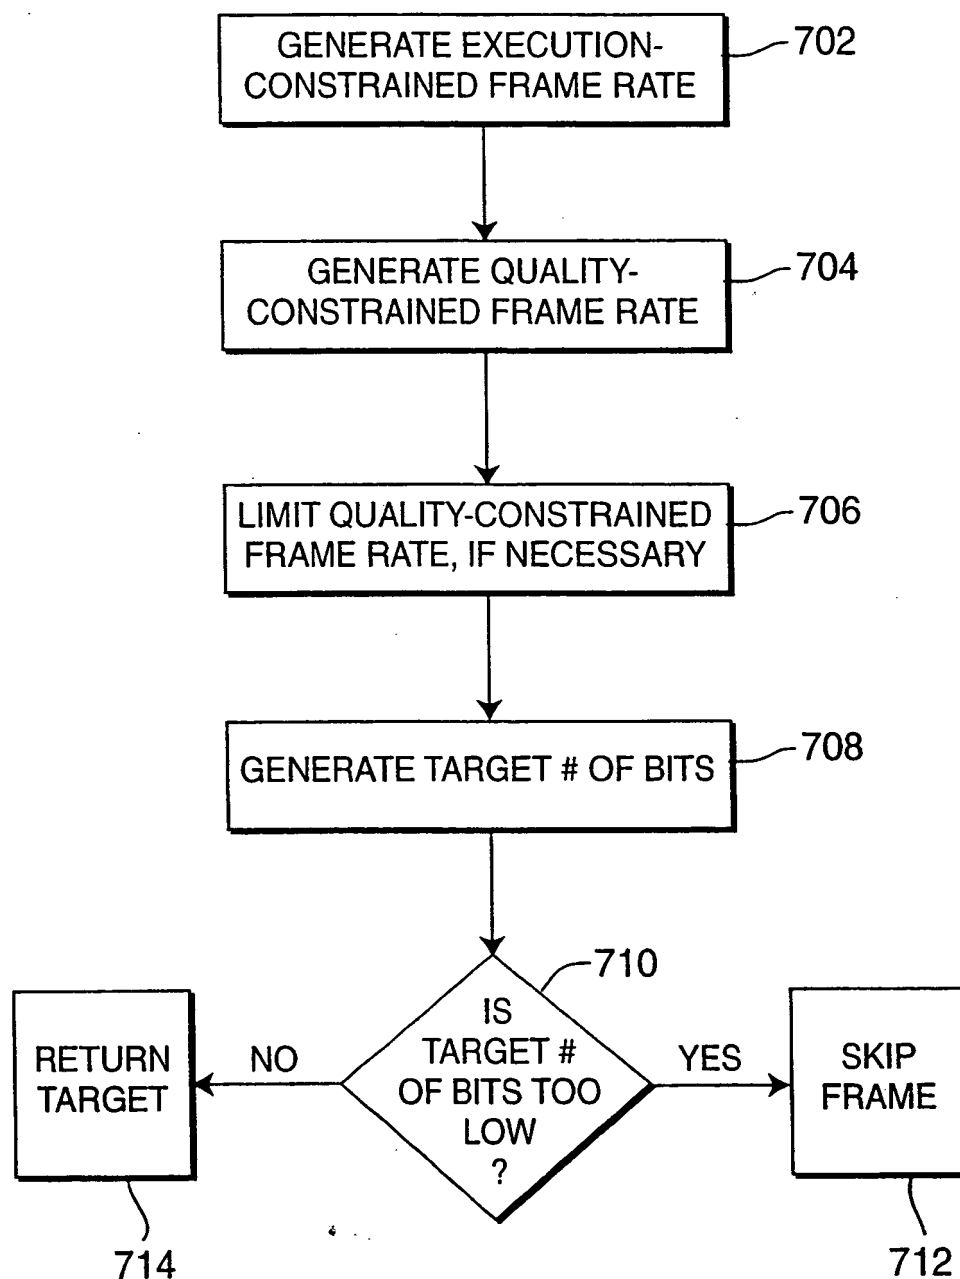
```
1 int rcFrameSkip (int bitrate, int skip_time_in_ms, int prev_bitCount)
2 {
3     Rslack += prev_bitCount - (prev_bitrate*skip_time_in_ms) /1000;
4
5     /* In slack calculations, the instantaneous bit rate at the previous
6        encoded frame is held */
7     /* until the current encoded frame. */
8     /* Under allocation slack cannot grow beyond R_WINDOW worth of bits */
9     if (Rslack < - (R_WINDOW*bitrate) /1000) Rslack = - (R_WINDOW*bitrate)
10        /1000;
11
12     prev_bitrate = bitrate;
13
14     /* Update P_frame_time which is the average time taken to code a P-frame. */
15     if (not I_frame)
16         P_frame_time = (P_frame_time*3 + skip_time_in_ms) >>2;
17     if (Rslack > (bitrate*(MAX_BUFDELAY_CONST*P_frame_time)) /1000)
18         return 1;
19     else return 0;
20 }
```

FIG. 4

```
1 int rcGetTarget (int Quality, int bitrate, int MIN_BITS_REQUIRED_FOR_CODING)
2 {
3     float frame_rate, target_frame_rate;
4     prev_bitrate = bitrate;
5     frame_rate = 1000.0/P_frame_time; /* Execution-constrained frame rate. */
6     target_frame_rate = def_frame_rate*(1.0 - (Quality-50) /P_SENSITIVITY);
7     /* Quality-constrained frame rate. */
8     /* Quality based frame_rate is possible only if execution frame rate allows it. */
9     if (target_frame_rate > frame_rate) target_frame_rate = frame_rate;
10    target_bits = (bitrate/target_frame_rate) - ((Rslack*P_frame_time) /
11    (R_WINDOW));
12    /* If target bits are too low, skip the frame. */
13    if (target_bits < MIN_BITS_REQUIRED_FOR_CODING) return -1;
14    else return target_bits;
15 }
```

FIG. 6

6/6

**FIG. 7**

INTERNATIONAL SEARCH REPORT

Int'l Application No

PCT/US 00/02800

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 H04N7/26 H04N7/50 H04N7/32

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04N

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
E, L	WO 00 18137 A (SARNOFF CORP) 30 March 2000 (2000-03-30) the whole document	1, 2, 4-6, 13, 17, 18
X	WO 98 37701 A (SARNOFF CORP) 27 August 1998 (1998-08-27) page 7, line 18 - page 10, line 35 page 12, line 19 - page 13, line 4 page 15, line 7 - line 31 page 17, line 3 - page 18, line 4	1, 2, 4-6, 10, 11, 13, 17, 18
Y A	US 5 333 012 A (SINGHAL SHARAD ET AL) 26 July 1994 (1994-07-26) column 3, line 29 - line 32 column 5, line 56 - line 64 column 7, line 53 - column 8, line 46	3, 14-16 9, 12
Y		3
	-/--	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

4 July 2000

Date of mailing of the international search report

17/07/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Ogor, M

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/02800

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5 852 669 A (JACQUIN ARNAUD ERIC ET AL) 22 December 1998 (1998-12-22) column 10, line 40 -column 11, line 8 column 11, line 50 -column 12, line 41 ---	15,16
Y	EP 0 836 329 A (SONY CORP) 15 April 1998 (1998-04-15) column 11, line 12 - line 26 column 11, line 33 - line 54 column 14, line 14 - line 16 ---	14
A		15,16
X	US 5 416 521 A (CHUJOH TAKESHI ET AL) 16 May 1995 (1995-05-16) column 4, line 39 -column 5, line 3 column 5, line 34 - line 43 column 6, line 8 - line 13 column 6, line 47 -column 7, line 15 column 9, line 10 - line 18 ---	1,2,4, 17,18
A		3,5,6, 10,11, 14,15
X	VETRO A ET AL: "Joint shape and texture rate control for MPEG-4 encoders" ISCAS '98 PROCEEDINGS OF THE 1998 IEEE INTERNATIONAL SYMPOSIUM ON CIRCUITS AND SYSTEMS, vol. 5, 31 May 1998 (1998-05-31) - 3 June 1998 (1998-06-03), pages 285-288, XP002141770 Monterey, CA, USA the whole document -----	1,2,4, 15,17,18
A		5,6,10, 11

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 00/02800

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 0018137	A	30-03-2000	JP 2000102008 A	07-04-2000
			WO 0018136 A	30-03-2000
			WO 0018134 A	30-03-2000
			WO 0018130 A	30-03-2000
			WO 0018135 A	30-03-2000
			WO 0018131 A	30-03-2000
WO 9837701	A	27-08-1998	CN 1247670 T	15-03-2000
			EP 0960532 A	01-12-1999
US 5333012	A	26-07-1994	JP 5167998 A	02-07-1993
US 5852669	A	22-12-1998	US 5500673 A	19-03-1996
			US 5512939 A	30-04-1996
			CA 2177866 A	11-01-1997
			EP 0753969 A	15-01-1997
			JP 9035069 A	07-02-1997
			CN 1118961 A	20-03-1996
			EP 0676899 A	11-10-1995
			US 5550580 A	27-08-1996
			US 5550581 A	27-08-1996
			US 5548322 A	20-08-1996
			US 5596362 A	21-01-1997
EP 0836329	A	15-04-1998	WO 9739588 A	23-10-1997
US 5416521	A	16-05-1995	JP 5336511 A	17-12-1993

THIS PAGE BLANK (USPTO)